

An Interview with
HOWARD FRANK

OH 188

Conducted by Judy O'Neill

on

30 March 1990

Fairfax, VA

Howard Frank Interview
30 March 1990

Howard Frank Interview
30 March 1990

Tape Index

Tape 1/Side 1.....	1
Side 2.....	17

Abstract

The interview begins with a discussion of Frank's background in networking, including his education, work in the Office of Emergency Preparedness, and founding of the Network Analysis Corporation. Frank then describes his work on the ARPANET, including his interaction with Roberts and the IPT Office, relationship with BBN and Leonard Kleinrock, and work on other network-related projects. The interview concludes with some general comments about IPTO as a government agency.

HOWARD FRANK INTERVIEW

DATE: 30 March 1990

INTERVIEWER: Judy

O'Neill

LOCATION: Fairfax, VA

O'NEILL: I'd like to start off with having you go over your background and experience before getting involved in the ARPANET project.

FRANK: Okay. Easy to do, since I became involved in the ARPANET project so early I didn't have much of a background before then. I got my Ph.D. in electrical engineering in 1965. I was on the faculty at Berkeley as first an assistant then an associate professor of electrical engineering and computer sciences at the University of California. I got a call one day in 1968 from the White House executive office of the President of the United States, asking me to be a consultant. I ended up spending a year on leave from Berkeley - the second half 1968 and the first half of 1969 - working in the Office of Emergency Preparedness, running their network analysis group. I met Larry Roberts at that time, in a casual way. I think that he was trying to sell me something, and I wasn't buying at the time. (Laugh) He was trying to sell me the idea of funding the thing, and I didn't have any money so I couldn't be buying anything. I think we actually had dinner once, and maybe I met him one other time. I think my view at that time was, "Who is this crazy guy?" to be honest.

O'NEILL: So he was at ARPA at the time?

FRANK: He was at ARPA at the time. He was just starting off with the idea of packet switching, etc. I had gone there for a one-year leave of absence from Berkeley and never went back to Berkeley, because after my leave I decided with another colleague of mine to start a company. Then I ran into Larry Roberts. My specialty was network analysis and he requested

a proposal from us regarding the specialty that we had, which was in the network analysis and design field. Essentially my work was taking the theoretical, sort of abstract, notions about how you design networks and applying them to the real world, and we had been quite successful. We had some really truly demonstrated successes. His idea was maybe we could do the same thing for the ARPANET and the whole packet switching technology in general. We

wrote a proposal, and we were awarded a very small contract in
→ October of 1969. I think we started working in October of 1969 to design the network structure of the ARPANET, to

redesign it. At that point in time, ARPANET was a four-node network. It existed on the west

coast. Larry had ordered maybe another twelve or fifteen lines across the country, and the

very first assignment was to put on a crash two-month effort, to see what we could do with the configuration that he had ordered. First of all, to analyze it because he had no tools for

analysis at all and second of all to figure out if we could find a better topology for it. I

actually remember the results very clearly. It ended up that we found the configuration that

would be thirty percent cheaper than the one that he had actually ordered and had forty

percent better performance in the way of through-put, simultaneously. He said, "That's great.

Let's do some more of it." (Laughter) This was a very little contract, I think it might have been

\$39,000, maximum, maybe it was fifty, but it was in that order of scheme of things. And from

that point on, since our specialty really was understanding networks, understanding analysis

and design, we began more and more work on the ARPANET, which was one of the early

projects that my first company, Network Analysis Corporation, did have.

O'NEILL: I'd like to go back to the year you spent in the Office of Emergency Preparedness.

What kind of network analysis were you doing at that point?

FRANK: We were part of something called the Systems Analysis Center. We were part of a little experimental group that had been the idea of Vice-president Humphrey, which was to see if analytic planning could work in a White House environment. You are too young to remember, but at that point in the time of Robert McNamara, remember the term the McNamara Whiz Kids in the Pentagon, were all the rage. And the White House had no comparable capability whatsoever. So the idea was that you bring some smart people in who could literally do something useful for government, in a White House environment. I got called in based on some people who heard me speaking. And they asked us to prove that we could do something - literally. That was the charter. To prove it in the context of the office of Emergency Preparedness. So we started working on two different classes of problems: one, network vulnerability and survivability, survivability of the country's assets in the event of a war, and second, economic analysis. And we picked a project with the Federal Power Commission that had to do with design of off-shore natural gas pipeline systems. That's a big mouthful of words. I had heard the word pipeline and I had seen pipes maybe on TV or certainly in the movies. I had never actually visited an oil field or anything like that. We came up with an actual computer technology and a computer program that saves \$100,000,000 in 1969. And Business Week wrote an article about it, and there were two hundred newspaper articles. It was a big blockbuster type thing, a very successful project of which the technology was given to the public. It was in the public domain; it was a government project. It wasn't the goal to invent pipeline technology but to prove that a group of smart people working in that environment could do something useful.

O'NEILL: Okay, so it was a proof of the organizational structure.

FRANK: Right, and the organizational strategy. Administrations changed, 1968 was an election year, Nixon came in. What's that old expression, "The operation was a success, but the victim died"? Well, that was exactly what happened because not only did they not agree to go forward on the concept of the thing, but in fact later on in time the Office of Emergency Preparedness was moved outside of the White House Environment. At the time that I was there, though, the head of the Office of Emergency Preparedness was Price Daniel who was the former governor of Texas and Lyndon Johnson's closest advisor, or one of his closest advisors. So that it was a strong organization with a goal in mind, and of course it changed, and they dropped the whole idea.

O'NEILL: How large of a group was it, do you recall?

FRANK: My group was, I had literally three or four full-time people, two of whom were my former students and two of whom were programmers that we borrowed from an agency. And I had about half a dozen consultants.

O'NEILL: Had you been working on the networking issues or operations? What kind of work had you been doing at Berkeley?

FRANK: My Ph.D. thesis was called "On Probabilistic Graphs and Some Applications." It was looking at two kinds of things: network theory as applied to communications networks and also network theory as applied to the connection of computers - the computers themselves - the insides of computers.

O'NEILL: How had you first gotten interested in that area, in communications networks?

FRANK: Somehow or other, when I was a senior in undergraduate school, I got interested in a course called network analysis, which had nothing to do with the kinds of ways we think about networks now. But it was circuit theory of how you connect up things that are wires and resistors and capacitors and vacuum tubes and some transistors and stuff like that. I kind of fell in love with that. And sort of my favorite thing when I was in high school was plane geometry. I never liked algebra, but I loved geometry because I'm a visually oriented person. When I went into graduate school, I was debating between becoming either a solid state electronics engineer, or a circuit theorist. And by luck, the professor who walked in, teaching the course in circuit theory, became a friend of mine. He was about ten years older than myself, and he was a specialist in graph theory, which is an abstract form of plane geometry - an elegant piece of mathematics. I fell in love with it, so instead of going into solid state electronics, I moved toward the circuit theory point of view. My master's thesis was called "Analysis of Networks by Vertex Reduction." Then I guess a year and a quarter after I started at Northwestern University I walked into my advisor's office one day and he was saying "I've been thinking," he just put his hand behind his head, and he said, "What if you took this network," and he draws a picture on the board, "and instead of links being there and certain, what if you didn't know too much about them and the capacities were random variables, as opposed to discrete deterministic events. How would you analyze that?" So that was an interesting problem.

O'NEILL: Did he have a particular application in mind when he was talking to you?

FRANK: He was thinking of communication networks where you're trying to pump things through the channels, and the channels are not necessarily there with fixed capacities, but they could randomly change because of either reliability problems, or maybe you don't know what they are. Therefore, you have some probabilistic guess about it. And he just formulated that problem. And I went out and I figured out how to solve it. And we wrote a letter to the editor of one of the circuit theory magazines on the thing. And I started extrapolating from there, and that's how my Ph.D. thesis grew up. I did it in minimum time. I went from bachelor's degree to Ph.D. in three years at Northwestern. I fell in love with the whole idea of networks and graph theory and all that. About half my course work was in mathematics, because I had to go deeply into the probability and statistics theory in order to be able to deal with the problems of randomly surviving and randomly existing networks. The other half was in the network end of the business. The usual electrical engineering curricula. I had intended to go into the space industry, but again my advisor said, "You know, I met some people at Berkeley, the chairman of the department. I think you might like it out there." So I went out to Berkeley. I had no intention of becoming a university professor. I fell in love with the place, and they happened to like me also, so I hired on. Hired on, that's easy to say - I went through twenty-five interviews I think. (Laugh) And the day before I was supposed to start work, the Board of Regents confirmed the position. Until then I didn't have a job. So I showed up in August of 1965, as a faculty member. And I continued that whole branch of research, and that was my primary branch of research. I did some very interesting stuff. One of the pieces of work I did actually won a prize in 1969 for the best paper of the year in the IEEE Communications Society Magazine. It was good stuff.

O'NEILL: Did you have any connection to the network experiment that was going on at the University of California?

FRANK: Only in a roundabout way. When I was finishing my Ph.D thesis, a guy walked into my office at Northwestern and said, "I just saw your thesis in the library." And my reaction was "My God!" and I ran over to the library, and there is a book called "Communication Networks: Stochastic Message Flow and Delay," by Len Kleinrock. And I opened it up and I said, "Thank goodness, it's a completely different approach." Same area, but a completely different view of the model, etc. that are in it. It really wasn't my Ph.D. thesis at all. Then when I was at Berkeley teaching I invited Kleinrock to come and be special guest lecturer, and Len and myself became friends and we've been friends now for, I guess, it's twenty-five years. The first time I met him it was in a summer extension course, it might have been in 1966. And he had been camping with his wife, and he came down to the airport and the car went out of control because the brakes got lost on the mountain road down. And he nearly got killed coming to my course. Then he took out the slides, glass slides they were using - I don't know if you've every seen 2x2 glass slides instead of the plastic stuff we use now - and they were all cracked. The car had crashed. He went down the side of the road banging on the mountain, trying to stop his car. After that we became, I would say long - term friends, although we didn't see each other that often, if you know what I mean. And we would share not only one another's ideas, but work, we would lecture in one another's courses.

O'NEILL: You mentioned that you met Roberts while you were in Washington. Why was he asking you for money?

FRANK: Because he was asking everybody for money, okay? It may have been Len Kleinrock who introduced me to Larry Roberts, it's very possible. Because Len, remember I said I had four full-time employees and some consultants? Len was one of my consultants at the White House. I had Len, I had a couple of people from Princeton, and I had a fellow from MIT (all very bright people) consulting. And I had a couple of my colleagues from Berkeley, all consulting on these problems. It probably was the case, because Len did know Larry Roberts. It probably was the case that Len introduced me to Roberts. You know, you say White House, and you write your name on White House stationery, and everybody thinks you've got more power than God. (Laugh) Right? Then power and money are equivalent, so they think you've got more money than God -and Larry was looking for money, period.

O'NEILL: Okay. Can you recall how far along the ARPA networking plan was?

FRANK: It was just a concept. It was really just a concept. It was the middle, or late 1968, so it was more than a concept, but to tell you the truth, I really wasn't paying that much attention. It was just an idea. My own work had been in a different end of the networking part of the business. It seemed like an interesting idea, but I was wrapped up in the things that I was doing, which were much more important. If I focused on it, it was not for very long, and if I did think about it in any great depth, I came up with the wrong conclusion, because I didn't think much of it, but that is, I think, because I didn't think much about it.

O'NEILL: You mentioned that you went to the Network Analysis Corporation at the end of 1969. Was it already an existing corporation?

FRANK: The way Network Analysis was formed was I took over a group of three or four employees at the OEP from another fellow who had been there for a year on leave. I directed the project that saved the \$100,000,000 of the Federal Power Commission.

O'NEILL: This is at...

FRANK: ... at the OEP. And he finished up his year's leave. As we were doing this thing... The way companies start is somebody says - we were saving a lot of money - and we said "If only we had a couple of percent of that," - that was a time when a million dollars was a lot of money - "we would be very happy." So I said, "Why don't we do it for ourselves?" And so he went out, he became a professor at Brooklyn Polytechnic.

O'NEILL: Who is this?

FRANK: A fellow named Dave Rosenbaum. I haven't seen him for a decade. He became a professor at Brooklyn Poly. He actually incorporated Network Analysis Corporation, but we did nothing until I finished my year, and I was the first employee. I became president of the company a year later, actually. He was only there for a year. He was going to be the business manager and the outside guy, and I was going to run the technical projects. And it didn't work out; he didn't work out and I took over, and I became the president and built the company up.

O'NEILL: So the experience that you had, that Roberts called upon, was your previous experience at the OEP as opposed to any experience at the Network Analysis Corporation?

FRANK: Yes. Network Analysis Corporation was just getting going. It was OEP plus my academic research, which was exactly in that topic. For instance, the prize winning paper that

I gave was called "Vulnerability of Communications Networks." It was a neat paper. That was why I ended up being invited to the White House in the first place, because I discovered a new way to analyze networks that people had been trying to figure out for years and spending millions of dollars of computer time - and I figured out how to do it with an equation. It was neat kind of stuff.

O'NEILL: Was that connected to the work Paul Baran and others had done at Rand in survivable networks?

FRANK: It was the follow-on to the work that Paul Baran did. Paul had the idea of distributed networks, and reliability indices, and looking at networks with degrees of density in the nodes, and even burned all kinds of computer time running graphs. What I did was I came up a theory of how to analyze those things, and I could reproduce Paul's graph with a paper and pencil and a calculator instead of computer time. I got the idea from work that I read about in the Journal of Mathematical Biophysics which was very interesting. I was just rummaging through, because I was looking for ideas, and I read a paper on a mathematical analysis of cell damage due to radiation. These guys were modeling cell damage, and there's particles coming in and bombarding, and I said, "Bombarding, that sounds right." (Laugh) It was risky work for Berkeley because it was during the Vietnam War, and the word bombing was not a particularly pleasant experience. But some people who became familiar with my work invited me out to Washington to give a talk at the Institute for Defense Analysis. Somebody from the White House was sitting in the audience and heard that work, you know that paper, which was really a very good thing. That's an objective statement, because

somebody else said it also, they gave me a certificate and \$300 which seemed like a lot of money at the time, for the paper. Then I got invited and I came out to Washington three months later, and within two hours they invited me to come spend a year.

3 O'NEILL: When you got involved, the network already existed with four nodes. Did you have the role in choosing where other nodes would be added?

FRANK: No. They would come and they say, "Here's where we want to put things." Lenny, Len Kleinrock, would come and say, "Here's our estimate of the traffic." And we would say, "Here's how you should connect them. Here's how they will perform." If you remember at the time, the ARPANET was being pushed for its reliability characteristics. Yet, nobody knew how to analyze the reliability of the things, and that had been my academic specialty. So I and another colleague of mine worked out a pretty interesting theory that we published in two or three papers on a better way to analyze the reliability of these networks. We built a model to analyze distributed reliability as opposed to through-put and delay. That was one model we had. We published that first paper. I think it was called "On Reliability Analysis I", in about 1972. Our work started expanding quite rapidly about that time because Larry would come up and have all these questions. He would go to two or three places, like the network measuring and monitoring center at UCLA and then for analytical-type stuff, or modeling, or topological kinds of things, he would come to us. Then around 1971 or so he asked us, "What would happen if you would take this stuff and you would expand it to very large networks?" I guess ARPANET was about fifteen nodes at the time, or planned at fifteen, or somewhere in there. He said, "Okay, what would happen if you put it in all the metropolitan centers in the United

States? How would you design it? What are its performance and curves codes, and all that?"

So we built models of that, computer models, which were also pushing state-of-the-art of computer modeling, and network modeling in particular. And the curves that you may or may not have seen that show cost per bit as a function of size, or cost per packet as a function of size, there were curves of number of nodes vs. cost per capacity unit. We generated all those, by actually simulating the construction of networks, going into population centers, generating traffic matrices, and then actually designing the nets and figuring out what they would cost to do that. It was a fairly sophisticated, very elaborate procedure.

O'NEILL: Did you need a lot of computing capacity in order to do these models?

FRANK: Yes, and we were pushing the state-of-the-art of that. We were using the Control Data 6600 machines to do that.

O'NEILL: Were you getting money from DARPA in order to have the 6600s?

FRANK: We didn't own the 6600s, we bought time from Control Data itself. It was a very interesting tradeoff because we didn't own the machines and because computing time was expensive and because every dollar (my budgets were fixed, literally) I spent on computer time was a dollar less I had for labor. Therefore we spent a tremendous amount of effort in trying to build computationally efficient models. And in doing so, and this was actually fundamental research over a time, we developed a whole family of computationally efficient algorithms. And we were publishing, we published dozens of papers all the time, all through this. And that was a stimulus to another whole field. You know, one of my guys developed a whole new

technology for efficient layouts of nets. And another guy came up with a better way to analyze the things, and we put all this stuff together. Each time always fighting the fact that every time we ran the 6600 it cost us a lot of money. My budgets were not hundreds and hundreds of thousands of dollars. I think at that time the largest funding I ever got from ARPA was about \$200,000 to work on the ARPANET, plus additional modeling to see how this would extrapolate into really large networks.

O'NEILL: Okay, so some of it was for research beyond the ARPANET in a sense?

FRANK: It was looking at how ARPANET technology would apply into the large scale world. For example, one of the things we were asked to do, I think this is around 1973, was to work with the Defense Communications Agency to see how, if you took all of the separate defense communications data networks, and you laid them out, and you put them into an integrated network using ARPANET technology, what the economics would look like. And we built a very extensive model of requirements and locations, and we literally took thirty different networks and redesigned each one of them to get a baseline. You can't compare it to the existing ones, because what if they were bad designs? So we redesigned each one to find the lowest possible cost free standing network. Then we created a family of designs that integrated everything through what is now called the Defense Data Network. That was actually one of the input studies to the funding of what originally was called AUTODIN II and is now called the Defense Data Network. That was also on DARPA funding. It was not specifically associated with ARPANET, although it was the extrapolation of ARPANET into the defense environment. Those were the kinds of things we did.

O'NEILL: How many people were working with you?

FRANK: I was the principal investigator. I guess I had at one time a little bit later, not on ARPANET but on some other things, I had as many as twenty people devoted to ARPA projects. In that timeframe, I guess it was maybe ten, not necessarily all full-time.

O'NEILL: Were you doing contracts with other government agencies at the same time? Or other companies?

FRANK: Actually DARPA was my only government contract at that time. Most of my other work was with commercial businesses. So for instance, at the same time we were doing ARPANET we were also testing and designing the Over-the-Counter securities network for the NASD. We actually had the assignment of field testing that thing for the NASD when it went into operation. It actually started about the same time as our ARPA contract. We would do things like that. We were also active in the cable television end of the business.

O'NEILL: Was the experience on the ARPANET helping in these other areas?

FRANK: It was helping because it was paying for part of the overhead. That's one real fact of life. But it was also helping because it was the more intellectually challenging of the work. And I could have really smart, capable people working on that, and also applying part of the time - because they weren't working full-time, I didn't have ten people working full-time on the ARPA contract. Clearly if you did that with \$200,000 a year, including overhead, they are not going to go that far. But I could take smart, really smart, intelligent guys, some of my colleagues from Berkeley, and apply them developing future models of how you do computer

network design, and at the same time use them in more mundane commercial applications, where, if you only put them into commercial applications, after a while they would just leave. They would get very burned out. So that was a specific effect that is hard to actually calibrate or measure. Later on, if I step back five years later on in the process, the work with ARPA helped a lot in a very direct way. We went out in the late 1970s and took a lot of the technology that we had developed for designing networks themselves, and a lot of the work that was done was originally done, or the ideas that were published, originally came out of some of the research that we did in designing ARPANET, and we built a family of network design tools that became a great commercial success, the dominant commercial network design tools in the business.

O'NEILL: How long did you work with the ARPANET, or did you have that contract?

FRANK: About 1974 as I recall. We delivered the design programs that we developed to the Defense Communications Agency and trained them to take over the design. Maybe 1973, maybe 1975, but it's around that timeframe. By that time very little of our work was really ARPANET related. And our work with ARPA had moved into two other areas: one, packet radio, where we became one of the original contractors on the packet radio project; The second was packet voice. I, as principal investigator of my company, did a major study of looking at the economics of packet voice. It was the principal economic study in the 1970s that did that. It was a major two-year project that we had done. Larry Roberts was no longer there; Robert Kahn had already taken over. I think, I'm not sure, but I think Bob may have started the project off, although we may have been talking with Larry about it.

TAPE 1 - SIDE 2

O'NEILL: You mentioned that Larry Roberts would come and talk about the additional nodes he wanted to add and ask you for guidance on how to connect them. What was your working relationship with him? Was it fairly informal?

FRANK: It was fairly informal. We would have program reviews, in the same way that everybody else had program reviews. But at that time we were really exploring the nature of networks. I would come to Washington periodically; it was rarely that he came to our offices, which were in New York. I would come to Washington, and we would talk about what we had learned. So it was not the formal program review part of it, but the "Gee whiz" kinds of things. That was a very useful thing. I don't know if that kind of environment can still exist at ARPA. It did at that time.

O'NEILL: Were those meetings just with you and Roberts?

FRANK: Generally.

O'NEILL: Did you have any interconnection with the other contractors? Did you have any need to coordinate with BBN for instance?

FRANK: We did work with BBN. We worked with UCLA. Those were our two primary interfaces. For instance, I think roughly 1972, Bob Kahn who was at BBN at that time, Len Kleinrock, and myself sat down and wrote a paper on experiences in design and theory. It was an interesting paper. The way it was done was interesting, because we had really come about it from completely different directions and learned that we discovered the same truths - if there are such things in this business.

O'NEILL: As a principal investigator on an ARPA contract, were you involved with the other principal investigators; did you go to the principal investigator meetings?

FRANK: Yes, of course. Always.

O'NEILL: Was there a lot of interest in the network from the other PIs?

FRANK: There was a lot of interest. There were a large number of very smart people there. I remember the first meeting that we held was in New Orleans, and Larry had asked me to give a talk on the network design.

O'NEILL: This is 1970, then?

FRANK: It might have been, you know? It certainly seemed like that. You would have to figure it by looking at lists of where the meetings were, whether or not it was New Orleans or not. But I know that it was in New Orleans, and it was the first meeting that real interest was evoked, and I remember what it was I reported on, because it was the initial, "Here is what is going to be done. Here is what we came up with." It was the only time that there was an actual reference point. Larry had a bill that he was going to be paying, and it had a set of parameters associated with it, and we came in and it was a six-week crash effort. People were just worked around the clock for six weeks, because he had a timetable he had to keep - there was a nine-month lead time on the lines. And six weeks later we said here's forty percent better performance and thirty percent less cost. And the technology we used was the same technology that we used in the design of the gas pipeline system, which was technology that grew out of something called "heuristic design for the traveling salesman problem." It is interesting because

we were theoreticians, but we were pragmatists to the extent that we said, "We don't know how to solve these problems in a purely mathematically optimal way." But somebody needs to get a result out, and you have to make a recommendation, and therefore we came up with the heuristic method of designing these networks that had to do with iterations, and essentially using repetitive analysis and smart goal-seeking types of programs. They could do better than a human being could do with output of the designs. Yet, it couldn't be proved to have any optimality in any sense, but it did better than.... It started off with what Larry finished with, and then it would improve on it, which is a neat kind of thing. About six, seven years later the theory of networking had advanced to the point where they proved the problems that we were solving were not solvable. If we had known that then, of course, we would have said the only way you can solve it is heuristics. We didn't know at that time because the theory didn't admit that kind of analysis, that we really didn't have a set of solvable problems anyway so you had to use heuristics.

O'NEILL: What kinds of reactions did you get to your talk at the PI meeting?

FRANK: Very good. Very good, because you see it was a neat, clean thing. "Here it is, here's what it is, here's a comparison." So I remember it generated a lot of enthusiasm, it was late in the evening, maybe 6:00, the last one, and it actually generated applause, which was unusual for a PI meeting. Those were fun meetings; those were very useful meetings.

O'NEILL: Can you give me an idea of how large this meeting was?

FRANK: Thirty people. Something like that.

O'NEILL: So, pretty small.

FRANK: Not a hundred and not five.

O'NEILL: How important was cost in optimizing the network designs?

FRANK: That was the driving criterion. The way we had established the goal seeking program, the objective was minimize costs with a set of constraints, where the constraints were on performance in terms of DARPA response time and reliability. So it would look to minimize costs subject to getting two tenths of a second response time and maintaining a certain level of reliability.

O'NEILL: Was it ever the case where the cost would be the limiting factor in the sense of you would degrade response time, for instance, in order to keep it under a certain cost?

FRANK: What we did was we did parameter analyses in order to see what would happen if we increased response time. The way these things work is they have one of the queuing curves and it makes no difference really. Originally Larry was talking half a second response time, if I recall. We did our analysis and it costs you the same to do a half a second as two tenths of a second because you are at the knee of the curve. So we lowered the constraints to two tenths of a second because it was better performance and it didn't cost anything more to get the two tenths of a second.

O'NEILL: Was there much debate about your findings? Were there times when Roberts said, "No, go reoptimize"?

FRANK: No, generally not.

O'NEILL: Generally, were you being a consultant to the office?

FRANK: We were the designers. We were the topological designers, and there were some times when Larry would say, "What if you did this?" It got to the point where it became a reliable enough process, that he would say, "What if I add these? What if I change the sequence in which I add these things, how would that work?" Or, "What if we change the traffic parameters?" But there was very little debate after that first time when we said, "Here it is." The thing is he could actually calculate the cost. As long as we were doing the arithmetic right, he could say ours is cheaper than theirs. After that there was never any "their's" anymore because ours was theirs. So I can't remember any controversy at all in the design part of it, because I don't know what he would have used or what anybody would have used as a base to contradict.

O'NEILL: How did you interact with BBN? It would seem that what you were doing, and what they were doing were closely related.

FRANK: It was somewhat of an adversarial relationship. Okay? It wasn't like a negative relationship, but it was always, "Here are a bunch of other smart guys." And we started meddling in things that they thought were exclusively in their domain. For instance, we would be starting... you can't just optimize topology, see that's impossible, because routing is a part of topology. So all of a sudden we started looking at routing algorithms, and they were saying, "Well, we are the routers." And we are saying, "Well, but you know if you did it this way..." (Laughter) ... "it would be better." There was one time where we said, "Well, if you do this, you get a better routing technique - period. It's just better." And I remember they said, "Well, you will have to work it out, because we don't believe it." And I had very smart guys, and they worked it out. (Laughter)

O'NEILL: So it wasn't a matter of having to have someone like Roberts arbitrate in that case, it was a matter of working it out until everyone was satisfied?

FRANK: Right. Arbitration wasn't needed because it really was a gentlemanly kind of thing. And maybe it was because we were both sensitive enough not to really crush one another's domains. There was always a little bit of discomfort there. Frank Heart was the project manager at the time, and Frank was - pardon me, he was the VP in charge of the project at the time. He had a very, "My team, my guys, my project, and I'll defend it with my life." You know, the mother wolf protecting her cubs. So that was just part of the thing, you know, part of the environment. It was an adversarial relationship, but it was not a hostile adversarial relationship, if that is consistent... if it makes sense.

O'NEILL: You mentioned before that you had done some work with the DoD, thinking about military applications for networks. When you were working on the ARPANET, were there ever any military considerations?

FRANK: Yes, because what happened was that ARPA was sensitive to that whole part of the thing. ARPA worried about that as well when we were looking at the application of the ARPANET, they had DCA working with us. So it was like a three - way team kind of thing. Then we were designing this technology for network design, and DCA gave us a contract to help develop the same technology to put it into their environment. So it was a sort of tri-partite arrangement where, in particular, a little bit later on, you had DCA who were looking at the things we were doing, and in fact in one case was the reviewer of what we were doing, you had ARPA, and then you had us. And we also had, I think, by that time a couple of other defense

projects. Not big deal type things, but I know we were working on something with the Air Force Office of the Secretary on survivability of military communication systems. And it was enough to be kind of aware of what was going on, not necessarily to have a really sharpened insight into, taking ARPANET and making it a complete defense network. But when we looked at the application of ARPANET to technology, to the defense environment, we had very much input from DCA, et cetera.

O'NEILL: So is it fair to say then, that your involvement with the military was not through ARPA so much as through your other contracts that dealt with your ARPANET experience?

FRANK: No. I think it was both, really, because ARPA introduced us to DCA and said, "We want you to look at these things, and we want you to use a model of, for instance, locations and traffic, that DCA will provide you." Okay, so ARPA was literally the contact point to DCA. In one of our reports they said, "Not only do we want you to look at the model, but DCA is going to be the reviewer of your work." And so when we prepared the report, the draft report was reviewed by DCA. And then we incorporated their comments as well as ARPA's draft comments.

O'NEILL: Did you get involved with the usage of the network?

FRANK: We became not a node on the network but a user of the network. Not for our first generation of design requirements, because we were crunching on Control Data. But in the second generation we developed a family of interactive tools using IMLAC technology. I don't know if you've every heard the expression. They were interactive graphics type of machines, about a \$30,000 machine that was very advanced. What we did was, if you looked at the

Control Data 6600, it was a batch technology and we ran a series of batch programs and a day later or two days later, you got a result. Then we came up with the idea of building an interactive family of programs where you had the man in the inner loop of the program. And you could only do that by using the ARPANET, and by using the DEC system tens that were on the network. It must have been around 1973, somewhere around there, we came out with our first version, because around that time the requirements were coming hot and heavy - check this thing, check that thing, what if you have ten nodes located over here. Because we needed to calibrate the programs, we ran it both ways, because I had one technology which was proven and working. And we were able to create the same design in a day that was taking us two weeks to get before that, using the actual ARPANET itself. And thereafter we disbanded the original set of tools, and the technology that we delivered to DCA to design the ARPANET was the technology that we built on top of the ARPANET.

O'NEILL: What were the difficult parts about dealing with the ARPANET? As it got larger, I would imagine it got increasingly more complex.

FRANK: Well, you see what we were doing was, as it got larger, we were way ahead of the actual size of the ARPANET. So that when it was a ten node network we were worrying about a thirty or forty node ARPANET. When it was a thirty or forty node network, we had already been asked to extrapolate this to the defense communications and the United States.

(interruption) So the hard thing was to sort of figure out how to design very large networks -

two hundred nodes, three hundred nodes. And we had to come up with a whole new technology for how to decompose network design work into pieces. And we did that. We did

the work there with two hundred or three hundred type node networks. We supported ARPA.
ARPA was trying to give away the ARPANET at one time, to get anybody to take it. And we generated a whole family of tradeoffs in analyses and designs that were presented not only to ARPA but also to AT&T. I and Larry Roberts went to a meeting at 195 Broadway - it has to be around 1973, around there I can't be sure, plus or minus a year, but I would guess it was 1973 - where Larry essentially made a presentation about the packet switch networks. I'm not sure we were calling it packet switching yet. That term seems to have popped up around 1972 or 1973. For instance, in 1970 we were not calling it packet switching.

O'NEILL: Okay. Do you remember what you were calling it?

FRANK: No. Computer networks, computer communications.

O'NEILL: But not...

FRANK: Store and forward. We were calling it store and forward networks. Store and forward computer networks. So somewhere around that 1973 timeframe packet switching pops up as a term, and Larry presented, at a formal presentation with slides, etc., on the economics, the values, the reliability, and all those things. And I presented the economic study that we had done on how this extrapolates into a major U.S. network. The results of the meeting were, if anybody was awake by the time we finished it, it was only because they were being courteous. And they said, "Go away."

O'NEILL: So it was not a negative reaction but a lack of interest?

FRANK: It was complete lack of interest, because they couldn't imagine why anybody would want to send data in the network. If I had to characterize the thing, I mean they were

telephone people - and that's what they were interested in. They had an opportunity, because the monopoly hadn't been broken yet, there was only one telephone company, to have picked up the technology essentially into the public domain for free. And they blew it. They really blew it.

O'NEILL: As far as you know there weren't any further attempts.

FRANK: As far as I know. I'm not sure what actually happened. I only participated in that single presentation. We published several papers, you know a lot of papers on the work of the ARPANET, but several on that itself. I don't remember where we published them. Somewhere I probably have one.

O'NEILL: Were there other people who were not interested or had a negative reaction to the kind of network you were talking about?

FRANK: I think it was just... nobody really cared. You know, it was earlier than its time.

O'NEILL: Let's end up with an open-ended question about the IPT office at ARPA and your involvement with them. Do you have any general comments you would like to make?

FRANK: Well, you know there was this evolution of IPTO at ARPA which had some very good managers, Larry Roberts followed by I guess Dave Russell. There may have been somebody in between, followed by Bob Kahn, maybe that's the sequence. Licklider was there; I guess maybe Licklider followed Larry.

O'NEILL: Right, for a short time.

FRANK: For a short time. And there was this open communication. It was government

operations, so you couldn't play any games or anything, but it was more of a creative environment that seemed to get lost sometime towards the late 1970s. I think the last work that we did was maybe in 1977, 1978 and we were no longer an ARPA contractor after that. Not for any special reason other than we sort of moved on in our direction, and ARPA kept moving in its direction. But it didn't seem to have the spontaneity that existed earlier on. I think at some point the principal investigator meetings stopped; I don't know why. Nobody asked me whether they should have them or not, but they stopped. Maybe they didn't stop and I wasn't invited, but I think they actually did stop. (Laugh) And those were very valuable kinds of things. Even though they weren't that frequent, they were quite valuable. It became more difficult to actually look for innovation. For instance in the management of our contract, it wasn't a very large contract; it was probably one of the smallest ones in the group because it was mostly thought/computer time as opposed to building a lot of hardware or whatever. But even in the context of that relatively small contract, there was room for somebody to ask the question, "Can you think about the following thing?" And we could go out and say, "Well, yes, that's part of our charter anyway, so let's just redirect this and think about the following thing," which is how we started thinking about the reliability of these things in a more formal way. And that seemed... it got outstructured, you know, that structure lost. On the other hand, if you look at the record of success, it was a fantastic office. And while I'm not very close to it now, from what I understand from, let's say reading about it as opposed to working with it, it's still really the place, and the government has to have a place. And I believe that military technology is the leading edge. I would expect that its future is still pretty sound. Although

probably not in the same way -- maybe you couldn't even have a little ARPANET come out again. Maybe it has to be a bigger program now. Although ARPANET was a big program in its time, it is little when you look in retrospect to a military platform or anything like that.

O'NEILL: You had had experience working in the government.

FRANK: Oh sure.

O'NEILL: Did the IPT office in particular strike you as different?

FRANK: Well you see the IPT office is five or six people. You see that's the thing. It's easy to say the "government," or ARPA, or something like that, but they are individuals that you deal with. And, you know, five or six people, the office is defined by those individuals. So indeed they were different, they were all skilled professionals, any one of them could have, did come from industry usually - picked themselves up and walked into higher paying jobs at higher levels, and I think they came back and forth with all the rest. So it wasn't a highly structured bureaucratic office. The philosophy was one of more of a stimulation of a group of people than telling them how to manage the work. So that's very different than, let's say, an FAA where it's a bureaucracy where you couldn't function with a group of independent professionals working towards a common mission. That's the way I would characterize IPT. In that point, we had a bunch of smart people who were working on a visionary basis with projects that sort of were coordinated, but they didn't have to be too tightly coordinated, which meant that the individual program managers could have a lot of influence over his particular program. It's not like fielding an SST and trying to make twelve billion dollars worth of equipment and people come together at a single point in time. You couldn't run a space

program like that, for instance.

O'NEILL: Okay, thank you.

END OF INTERVIEW